CMS Activities in FermiGrid

Ian Fisk

Current CMS Clusters

- CMS Currently has 4 Clusters
  - Grid2003 Enabled Cluster of 40 dual Athlon Nodes
  - 40 Dual Xeon 2.4GHz System Used for CMS prod
    - Planned to enable with Grid2003 interfaces
  - 80 Dual Xeon 3.0GHz Systems with Condor enabled. Plan to install both LCG and Grid2003 interfaces
  - UAF Interactive and User batch nodes: Currently 36 and ramping to 56 nodes, which are a combination of 2.4 and 3.0GHz Xeon systems

Why so many clusters?

- US-CMS has used new procurements to test implementations and new components
  - New Farm is being installed with Condor instead of adding it to the FBS resources
- The other reason is that scaling limitations were observed in the FBSNG job manager implementation in the GRAM bridge requiring us to partition our two production clusters
- We expect to reduce the number of production clusters as we get more Condor experience
  - 2005 resources should be added to the newest 2004 resources.

Planned Facilities

- At FNAL US-CMS is expecting to procure 300 new dual Xeon 3.X GHz nodes in FY2005
  - Enable with interoperable LCG/OSG interfaces

Opening CMS Systems

- US-CMS believes that allowing opportunistic use of US-CMS facilities can be done with reasonable effort
  - We would be interested in seeing the SAM-Grid job manager enabled on top of the Grid2003/OSG services on CMS facilities
    - For D0 applications we would anticipate them using D0 SAM Stations already at FNAL, but SAM client software will be installed for CMS testing as well
  - The CAF services from CDF might run more easily against the CMS Condor resources
    - CMS is interested in achieving the interactive/batch balance that CDF has seen in the CAF for our own analysis system

Grid Interfaces

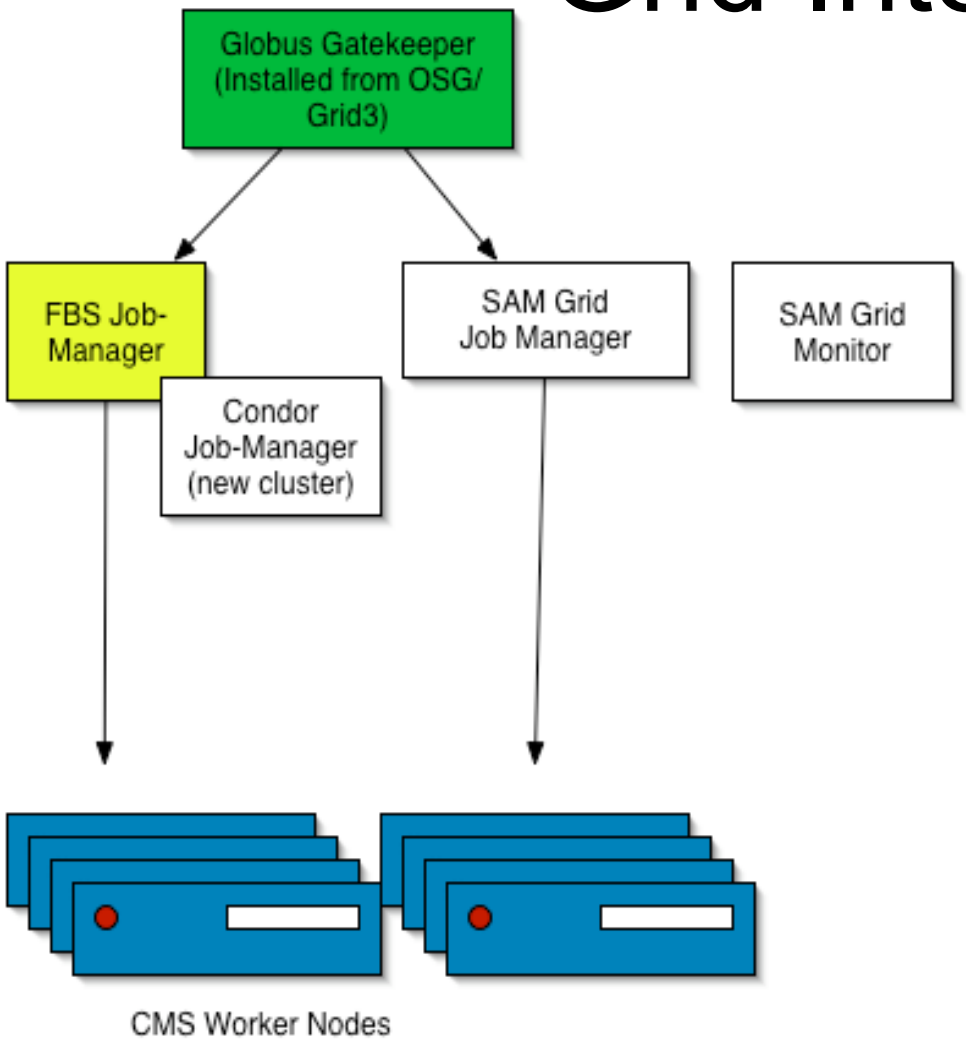

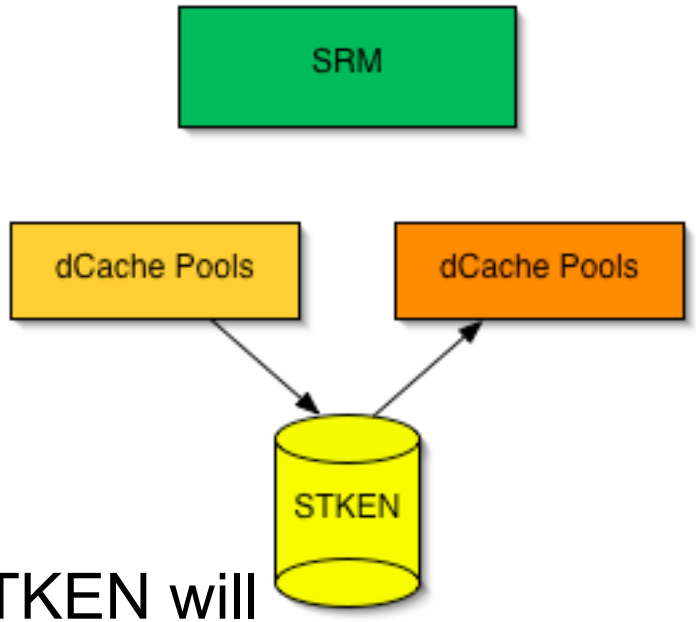

STKEN will Probably be used to support multiple VOs in OSG

- Items in white need to be added

Opportunistic Use

- CMS computing resources are busy
  - There is a significant increase in the expected MC output
- At the same time, the CMS critical periods are different from the running experiments
  - We don’t have big increases in activity before summer and winter conferences
- We do have structure in our usage patterns and there are opportunities for opportunistic use.

Support Load

- There is no intrinsic reason that supporting the grid interfaces to a cluster is any more operational load than supporting the ssh interfaces to a cluster
  - The software is upgraded with some frequency, but at least in Grid3 that has been once every couple of months
  - The packaging and configuration are improving
    - More work is probably needed
- We have seen some places the grid interfaces cause higher operational loads
  - In general these are easily identified as areas where functionality is missing
    - Local Storage fills up because storage management is missing